Noname manuscript No.
(will be inserted by the editor)

Metropolis-type algorithms for Continuous Time Bayesian Networks

Błażej Miasojedow · Wojciech Niemiro · John Noble · Krzysztof Opalski

arXiv:1403.4035v1 [stat.ME] 17 Mar 2014

Received: date / Accepted: date

Abstract We present a Metropolis-Hastings Markov chain Monte Carlo (MCMC) algorithm for detecting hidden variables in a *continuous time Bayesian network* (CTBN), which uses reversible jumps in the sense defined by Green (1995). In common with several Monte Carlo algorithms, one of the most recent and important by Rao and Teh (2013), our algorithm exploits *uniformization* techniques under which a continuous time Markov process can be represented as a *marked Poisson process*. We exploit this in a novel way. We show that our MCMC algorithm can be more efficient than those of likelihood weighting type, as in Nodelman et al. (2003) and Fan et al. (2010) and that our algorithm broadens the class of important examples that can be treated effectively.

Keywords Continuous time Bayesian networks · Markov chain Monte Carlo · Metropolis algorithms · uniformization

B. Miasojedow
Institute of Applied Mathematics, University of Warsaw
Banacha 2, 02-097 Warsaw, Poland
E-mail: bmia@mimuw.edu.pl

W. Niemiro
Institute of Applied Mathematics, University of Warsaw
Banacha 2, 02-097 Warsaw, Poland
E-mail: wniem@mimuw.edu.pl

J. Noble
Institute of Applied Mathematics, University of Warsaw
Banacha 2, 02-097 Warsaw, Poland
E-mail: noble@mimuw.edu.pl

K. Opalski
Institute of Applied Mathematics, University of Warsaw
Banacha 2, 02-097 Warsaw, Poland
E-mail: krzysztof.opalski@mimuw.edu.pl

1 Introduction

Continuous time Bayesian networks (CTBNs) represent explicitly temporal dynamics in probabilistic reasoning. They were introduced by Schweder (1970) under the name *Composable Markov Chains* and then reintroduced by Nodelman et al. (2002) as *Continuous Time Bayesian Networks*. A *continuous time Bayesian network* (CTBN) is a time homogeneous Markov process, which is decomposed into processes whose transition intensities depend on the other processes in the network. The dependence structure between the intensities is encoded by a graph. CTBNs are only loosely connected with Bayesian networks (BNs); the common feature is that of *modularity*, splitting a large problem into smaller components, where the relations between the components are represented graphically. The features represented by the respective graph structures are essentially different, although some of the techniques share similarities. For example, ideas from the *intervention calculus* of Pearl (1995) for BNs transfer quite naturally to CTBN analysis. CTBNs provides a promising and flexible class of models with applications in, for example, Survival Analysis.

While it is relatively straightforward to propose a CTBN as a statistical model, learning algorithms are computationally expensive. There are three categories of learning problem. The first is *answering queries*; that is, inserting evidence (information on the trajectory) for some nodes of the CTBN and making probabilistic inference about the behaviour of others when the structure and parameters are known. This is a classical problem in hidden Markov models (HMMs). The second problem is to estimate the parameters of a CTBN, that is the (conditional) transition intensities, based on a *learning sample*. The third and by far the most diffi-

cult problem is to learn the structure, that is the graph of a CTBN.

These tasks are 'hierarchical' in nature; approaches to the second and third tasks are based on and develop approaches to the first task. Even the first of the these tasks is challenging. There are deterministic algorithms, for example those proposed by the originators of the model, Nodelman et al. (2002) and developed further by Nodelman (2007). The performance of these algorithms has been examined through simulated case studies, but theoretical bounds on the accuracy of approximations have yet to be derived. Niemiro (2014) presents a *randomized* approximation scheme, which uses a deterministic algorithm for static BNs as a subroutine.

Remarkable progress has been achieved via Monte Carlo (MC) algorithms for CTBNs. Nodelman et al. (2003) proposed an algorithm which is a modification of the classical *likelihood weighting* (LW) method (Dagum and Luby (1997)). Recent results on this method are in Fan et al. (2010). The well-documented drawback of LW is the degeneracy of weights; the likelihood weights tend to be extremely unequally distributed, especially when there is a lot of data available. One remedy is to use *Markov chain Monte Carlo* (MCMC) techniques instead of independent sampling, while another is to use sequential Monte Carlo (SMC). An application of this method to CTBNs is given in Ng et al. (2005).

We consider MCMC algorithms based on the Metropolis–Hastings scheme. The idea originated in Metropolis et al. (1953) and was extended in Hastings (1970) and is at the core of almost all later MCMC developments. Our algorithm makes essential use of the representation of sample paths of a continuous time Markov process by a marked Poisson process (MPP). This representation, also known under the name of 'uniformization', has been exploited in many papers; one of the most recent is the work of Rao and Teh (2013). We use uniformization to construct a variety of proposals including reversible jumps, moves which change dimension of the space.

2 The basic algorithm

In this section we describe the basic algorithm in the general setting of continuous time Markov processes. In the following sections we use this algorithm as the main building block for more complicated procedures, specifically designed for CTBNs.

2.1 Continuous time Markov processes

We first recall some definitions and introduce notations which will be used throughout. Consider a continuous time stochastic process $(X(t), t \geq 0)$ defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ with a discrete state space $\mathcal{X}$. Generic states will be denoted by $x, x', x'', \ldots \in \mathcal{X}$. Assume that the process is a continuous time, *time homogeneous Markov process* (THMP) with transition probabilities

$$P^t(x, x') = \mathbb{P}\left[X(t+s) = x' | X(s) = x\right].$$

The initial distribution is denoted by $\nu(x) = \mathbb{P}\left[X(0) = x\right]$. As $\mathcal{X}$ is discrete, ν can be viewed as a vector and P^t as a matrix (both possibly infinite arrays). The *intensity matrix* is defined as

$$Q(x, x') = \lim_{t \to 0} \frac{1}{t} \left[P^t(x, x') - I(x, x')\right],$$

where $I = P^0$ is the identity matrix. Thus $Q = \frac{\mathrm{d}}{\mathrm{d}t} P^t|_{t=0}$.

Expressed differently, $Q(x, x')$ is the intensity of jumps from x to x':

$$\mathbb{P}\left[X(t + \mathrm{d}t) = x' | X(t) = x\right] = Q(x, x')\, \mathrm{d}t \quad \text{for } x \neq x';$$
$$\mathbb{P}\left[X(t + \mathrm{d}t) = x | X(t) = x\right] = 1 - Q(x)\, \mathrm{d}t,$$

where

$$Q(x) = -Q(x, x) = \sum_{x' \neq x} Q(x, x')$$

denotes the total intensity of jumping from x. Clearly, $\sum_{x'} Q(x, x') = 0$. The uniformization technique requires that $\sup_x Q(x)$ is bounded. We will assume that this assumption is satisfied.

Equivalently, a THMP satisfying this assumption can be described as a *marked Poisson process*. Consider the following sampling algorithm which uses *marking* and *thinning*. Let $\lambda \geq \max_x Q(x)$. At the first stage we sample *potential* moments of jumps, say $T_1 < \cdots < T_i < \cdots$. These are points of a time homogeneous Poisson process with intensity λ. At the second stage, we mark them. The marks, denoted $X_1, \ldots, X_i, \ldots$, are consecutive states of the *redundant skeleton Markov chain* with transition probabilities

$$P(x, x') = \mathbb{P}[X_i = x' | X_{i-1} = x] = \begin{cases} Q(x, x')/\lambda & \text{if } x \neq x'; \\ 1 - Q(x)/\lambda & \text{if } x = x'. \end{cases} \quad (1)$$

Now let $X(t) = X_{i-1}$ for $T_{i-1} \leq t < T_i$ $(i = 1, 2, \ldots,$ with $T_0 = 0$ and $X_0 \sim \nu)$. The process $(X(t), t \geq 0)$ is a THMP with intensity matrix Q and initial distribution ν. In the sequel we fix a finite time interval, say

$[t_{\min}, t_{\max}]$. In this section, to simplify notation we set $t_{\min} = 0$ and $t_{\max} = 1$. The process

$$\boldsymbol{\Xi} = (X(t),\ 0 \le t \le 1)$$

is thus represented by

$$\mathbf{X} = \begin{pmatrix} 0 & T_1 & \cdots & T_i & \cdots & T_N & 1 \\ X_0 & X_1 & \cdots & X_i & \cdots & X_N & \end{pmatrix},$$

where $N = \max\{n : T_n < 1\}$, with the corresponding sample path

$$\mathbf{x} = \begin{pmatrix} 0 & t_1 & \cdots & t_i & \cdots & t_n & 1 \\ x_0 & x_1 & \cdots & x_i & \cdots & x_n & \end{pmatrix}. \tag{2}$$

The construction described above is also known under the name *uniformization of a THMP*, for example Hobolth and Stone (2009) and Rao and Teh (2013). The representation (2) is *redundant* in the sense that there are infinitely many different double sequences $\mathbf{x}$ which correspond to the same sample path $\boldsymbol{\xi} = (x(t), 0 \le t \le 1)$ of the THMP. We distinguish between an actual sample path $\boldsymbol{\xi}$ and a representation $\mathbf{x}$, by using Greek and Latin letters, respectively. The chief advantage of uniformization is that the sequences $T_1, \ldots, T_i, \ldots$ and $X_1, \ldots, X_i, \ldots$ are *independent* of each other. In this section we will work with representation (2). Let us write the probability density of $\mathbf{X}$ in the following way:

$$\pi(\mathbf{x}) \propto \lambda^n \mathrm{d}t_1 \cdots \mathrm{d}t_n \mathbb{I}(0 < t_1 < \cdots < t_n < 1) \\ \nu(x_0) P(x_0, x_1) \cdots P(x_{n-1}, x_n). \tag{3}$$

This is a slight abuse of notation, since by $\pi(\mathbf{x})$ we really mean

$$\mathbb{P}\big[t_1 \le T_1 < t_1 + \mathrm{d}t_1, \ldots, t_n \le T_n < t_n + \mathrm{d}t_n, \\ X_0 = x_0, X_1 = x_1, \ldots, X_n = x_n\big].$$

2.2 Hidden Markov models

Let $(X(t), 0 \le t \le 1)$ be a THMP. Suppose that process $X(t)$ cannot be observed directly, but that we observe *evidence* y, which is a realisation of a random variable Y with probability distribution $L(.|\mathbf{x})$. The quantity $L(y|\mathbf{x})$ is the *likelihood* of evidence y given a trajectory $\mathbf{x}$. We assume that the likelihood only depends on $\mathbf{x}$ through the actual sample path $\boldsymbol{\xi} = (x(t), 0 \le t \le 1)$; it does not depend on the 'redundant representation'. In this section the concrete form of the evidence (e.g. the space in which Y takes values etc.) is irrelevant.

The problem is to estimate the hidden trajectory $(x(t),\ 0 \le t \le 1)$ given y. From the Bayesian perspective, the goal is to compute/approximate the posterior:

$$\pi(\mathbf{x}|y) \propto \pi(\mathbf{x}) L(y|\mathbf{x}).$$

The function L, the transition probabilities Q and the initial distribution ν are assumed to be known.

In principle, the problem can be solved by the method of *likelihood weighting* (LW), but an inherent problem of LW is the degeneracy of the weights. Even in a relatively easy example such as that presented in Section 5, the efficiency of LW is poor.

2.3 The Markov chain Monte Carlo algorithm

We propose a version of the Metropolis-Hastings algorithm (MHA) which converges to the target distribution $\pi(\mathbf{x}|y)$. The general scheme is standard. At each step of the algorithm, we proceed as follows. Suppose that $\mathbf{X}_{m-1} = \mathbf{x}$. We first sample a *proposal* $\mathbf{X}' = \mathbf{x}' \sim q(\mathbf{x}, \cdot)$. Then let $\mathbf{X}_m := \mathbf{x}'$ (accept the move from $\mathbf{x}$ to $\mathbf{x}'$) with probablility $a(\mathbf{x}, \mathbf{x}')$ or let $\mathbf{X}_m := \mathbf{x}$ (reject the move from $\mathbf{x}$ to $\mathbf{x}'$) with probablity $1 - a(\mathbf{x}, \mathbf{x}')$. The general MHA recipe for the acceptance probability is

$$a(\mathbf{x}, \mathbf{x}') = \min\left(\frac{\pi(\mathbf{x}')L(y|\mathbf{x}')q(\mathbf{x}', \mathbf{x})}{\pi(\mathbf{x})L(y|\mathbf{x})q(\mathbf{x}, \mathbf{x}')}, 1\right). \tag{4}$$

function *StepMH*($\mathbf{x}$)

```
Sample x' ~ q(x, ·) { proposal }
Sample U ~ Unif(0, 1)
if U < a(x, x') then
   return x' { move accepted with probability a(x, y) }
else
   return x { move rejected with probability 1 − a(x, y) }
end if
```

Representation (2) suggests several choices of proposal q. We describe some of them below.

Change of time

Let $\mathbf{x}$ be given by (2). Leaving n and $(x_i)_{i=0}^n$ unchanged, we sample $(t'_i)_{i=1}^n$ as follows: for a given pair of times (t_i, t_{i+2}), we replace t_{i+1} with $t'_{i+1} \sim \mathrm{Unif}(t_i, t_{i+2})$, or we can similarly replace several (or all) of the times. The pseudo-code below gives this more precisely. We adopt the convention that $t_0 = 0$ and $t_{n+1} = 1$.

function *ChangeTime*($\mathbf{x}$)

if $n > 0$ **then**
Choose a pair (i_1, i_2) from the set $\{0, \dots, n+1\}$ in such a way that $i_1 + 1 \leq i_2 - 1$.
Let $t'_i := t_i$ for $i \leq i_1$ and for $i \geq i_2$
Sample new times $t'_{i_1+1}, \dots, t'_{i_2-1} \sim Unif(t_{i_1}, t_{i_2})$
Sort t'_is
end if
return $\mathbf{x}'$

When this move is proposed, the acceptance probability (4) reduces to:

$$a(\mathbf{x}, \mathbf{x}') = \min\left(\frac{L(y|\mathbf{x}')}{L(y|\mathbf{x})}, 1\right). \tag{5}$$

This follows from (4) by noting, from (3), that $\pi(\mathbf{x}) = \pi(\mathbf{x}')$. The points $t_{i_1+1}, \dots, t_{i_2-1}$ are distributed as a sorted sample from Unif(t_{i_1}, t_{i_2}). This follows from elementary and well-known properties of the Poisson process. The moves $\mathbf{x} \to \mathbf{x}'$ and $\mathbf{x}' \to \mathbf{x}$, given the choice (i_1, i_2), both amount to choosing uniformly distributed sets of random times with the same number of elements over the same interval. It therefore follows that $q(\mathbf{x}, \mathbf{x}') = q(\mathbf{x}', \mathbf{x})$ and hence *ChangeTime* is π-reversible; $\pi(\mathbf{x})q(\mathbf{x}, \mathbf{x}') = \pi(\mathbf{x}')q(\mathbf{x}', \mathbf{x})$.

Note that no restriction has been place, so far, on the way that the pair (i_1, i_2) is seleted. There are several possibilities. One way to implement *ChangeTime* is to impose the constraint $i_1 + 1 = i_2 - 1$. This amounts to sampling a single point t'_i. Another extreme is to choose $i_1 = 0$, $i_2 = n + 1$ and sample all t'_is. The move then loses its 'local' character, but is still π-reversible.

Change of skeleton

Let $\mathbf{x}$ be given by (2). We leave n and $(t_i)_{i=1}^n$ unchanged and sample $(x'_i)_{i=1}^n$. As with 'Change of time', there are several possibilities, ranging from 'local' to 'increasingly global'. Below we describe one of these; the function *ChangeState* updates only one state of the skeleton and is actually a variant of the Gibbs Sampler for discrete time chains.

The value of i is chosen uniformly over $\{0, 1, \dots, n\}$. The *proposal* transition clearly defines a π-reversible chain (by construction) and the acceptance probability is given by (5).

If sampling from the Markov bridge is not feasible, then the sampling may be replaced by a Metropolis step which targets this distribution. This can happen, for example, when the state space is large.

function *ChangeState*($\mathbf{x}$)

if $n > 0$ **then**
Choose i from the set $\{0, 1, \dots, n\}$
Let $x'_j := x_j$ for $j \neq i$
if $i = 0$ **then**
Sample new state x'_0 with probability proportional to $\nu(x'_0)P(x'_0, x_1)$
else if $i = n$ **then**
Sample new state x'_n with probability proportional to $P(x_{n-1}, x'_n)$
else
{ if $0 < i < n$ }
Sample new state x'_i with probability proportional to $P(x_{i-1}, x'_i)P(x'_i, x_{i+1})$
end if
end if
return $\mathbf{x}'$

Change of dimension

We now proceed to proposals which change the number of marked Poisson points (t_i, x_i). We describe several alternative moves of this type. Their relative merits probably depend on the parameters of the process Q, ν and L. We restrict ourselves to the moves which increase / decrease the number of points by one; that is, to algorithms of the following general form:

Sample $d \in \{-1, 1\}$ u.a.r.
if $d = 1$ **then**
Attempt a move that adds one point
else if $n > 0$ **then**
Attempt a move that erases one point
else
Do nothing
end if

Below we describe pairs of moves: a move which creates a new point together with its counterpart which anihilates one point. To clarify the proofs of reversibility, $\mathbf{x}$ will always denote a configuration of n points given by (2), counting neither $t_0 = 0$ nor $t_{n+1} = 1$, while $\mathbf{x}'$ will be a configuration of $n - 1$ points. When x_i is deleted from $\mathbf{x} = (x_1, \dots, x_n)$ to obtain $\mathbf{x}' = (x'_1, \dots, x'_{n-1})$ or vice versa, then $\mathbf{x}'$ in terms of the elements of $\mathbf{x}$, $\mathbf{x}'$ is: $(x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n)$ and this notation will be used.

The proofs that the pairs of moves are reversible will be sketched; the reader is referred to Green (1995) for the additional analysis required for full proofs.

The moves in the following pair: *EraseRandomPoint* and *AddRandomPoint* are designed to act locally. Suppose a time $t_* \in (t_i, t_{i+1})$ is proposed. We would like the corresponding site $x_* \in \mathcal{X}$ to be sampled according

to the mechanism

$$\mathbb{P}(x_*) = \frac{P(x_i, x_*)P(x_*, x_{i+1})}{\sum_{y\in\mathcal{X}} P(x_i, y)P(y, x_{i+1})}.$$

This would be the natural 'Markov bridge' between the two existing points. The denominator, though, may be computationally expensive for large $\mathcal{X}$ and therefore the simpler proposal $\mathbb{P}(x_*) = P(x_i, x_*)$ is used.

function *EraseRandomPoint*($\mathbf{x}$)

Select a subscript $i \in \{1, \ldots, n\}$ uniformly at random { remove t_i and x_i from $\mathbf{x}$ }
for $j := 1$ to $i-1$ **do**
  $t'_j := t_j$; $x'_j := x_j$
end for
for $j := i$ to $n-1$ **do**
  $t'_j := t_{j+1}$; $x'_j := x_{j+1}$
end for
return $\mathbf{x}'$

function *AddRandomPoint*($\mathbf{x}'$)

Sample t_* from $Unif(0,1)$ { add t_* to $\mathbf{x}'$ }
Find i such that $t'_{i-1} < t_* < t'_i$
for $j := 1$ to $i-1$ **do**
  Let $t_j := t'_j$ and $x_j := x'_j$
end for
Sample x_* from $\mathcal{X}$ according to $\mathbb{P}(x_*) = P(x_{i-1}, x_*)$
Let $t_i := t_*$ and $x_i := x_*$
for $j := i$ to $n-1$ **do**
  $t_{j+1} := t'_j$; $x_{j+1} := x'_j$
end for
return $\mathbf{x}$

For these functions, the acceptance are as follows.

$$a(\mathbf{x}, \mathbf{x}') = \begin{cases} \min\left(\frac{n}{\lambda}\cdot\frac{L(y|\mathbf{x}')}{L(y|\mathbf{x})}\cdot\frac{P(x_{i-1},x_{i+1})}{P(x_i,x_{i+1})}, 1\right), & \text{if } i < n; \\ \min\left(\frac{n}{\lambda}\cdot\frac{L(y|\mathbf{x}')}{L(y|\mathbf{x})}, 1\right) & \text{if } i = n, \end{cases}$$

$$a(\mathbf{x}', \mathbf{x}) = \begin{cases} \min\left(\frac{\lambda}{n}\cdot\frac{L(y|\mathbf{x})}{L(y|\mathbf{x}')}\cdot\frac{P(x_i,x_{i+1})}{P(x_{i-1},x_{i+1})}, 1\right) & \text{if } i < n; \\ \min\left(\frac{\lambda}{n}\cdot\frac{L(y|\mathbf{x})}{L(y|\mathbf{x}')}, 1\right), & \text{if } i = n. \end{cases}$$

Different choices of proposal $\mathbb{P}(x_*)$ for *AddRandomPoint* will alter the acceptance rates.

Proof (Reversibility of Add/Erase Random Point) The proof is sketched; a complete proof can be constructed quite easily along the lines found in Green (1995). Let $\mathbf{x}'$ denote the point with $n-1$ states and $\mathbf{x}$ the point with n states. Consider $\mathbf{x}' \to \mathbf{x}$ via *AddRandomPoint* and $\mathbf{x} \to \mathbf{x}'$ via *EraseRandomPoint*. The formulae for $\pi(\mathbf{x})$ and $\pi(\mathbf{x}')$ are respectively:

$$\begin{aligned}\pi(\mathbf{x}) &= \lambda^n \mathrm{d}t_1 \cdots \mathrm{d}t_i \cdots \mathrm{d}t_n \nu(x_0) \cdots P(x_{i-1}, x_i) \\ &\qquad P(x_i, x_{i+1}) \cdots, \\ \pi(\mathbf{x}') &= \lambda^{n-1} \mathrm{d}t'_1 \cdots \mathrm{d}t'_{i-1} \cdots \mathrm{d}t'_{n-1} \nu(x'_0) \cdots \\ &\qquad P(x'_{i-1}, x'_i) \cdots \\ &= \lambda^{n-1} \mathrm{d}t_1 \cdots \mathrm{d}t_{i-1} t_{i+1} \cdots \mathrm{d}t_n \nu(x_0) \cdots \\ &\qquad P(x_{i-1}, x_{i+1}) \cdots\end{aligned}$$

while the proposals are:

$$q(\mathbf{x}, \mathbf{x}') = \frac{1}{n}, \quad q(\mathbf{x}', \mathbf{x}) \quad = \mathrm{d}t_i P(x_{i-1}, x_i).$$

Therefore, for $i < n$,

$$\frac{\pi(\mathbf{x}')L(y|\mathbf{x}')q(\mathbf{x}', \mathbf{x})}{\pi(\mathbf{x})L(y|\mathbf{x})q(\mathbf{x}, \mathbf{x}')} = \frac{n}{\lambda}\cdot\frac{L(y|\mathbf{x}')}{L(y|\mathbf{x})}\cdot\frac{P(x_{i-1}, x_{i+1})}{P(x_i, x_{i+1})}$$

and the conclusion follows; similarly for $i = n$.

Finally, we propose a pair of moves which is essentially a restricted version of the previous pair, where we only allow a change at a *virtual* jump point. These moves are *EraseVirtualPoint* and *AddVirtualPoint*. The first of these removes a point x_i from the skeleton if and only if it is a *virtual point*, i.e. if and only if $x_{i-1} = x_i$. The second of these adds a virtual point into the skeleton. These may be used instead of *EraseRandomPoint/AddRandomPoint*, provided λ is chosen so that $1 - \frac{1}{\lambda}\max_x Q(x)$ is sufficiently large; if $\max(1 - \frac{1}{\lambda}Q(x))$ is close to zero, then there will be few virtual points; the *EraseVirtualPoint* algorithm will rarely find a virtual point and the *AddVirtualPoint* algorithm will reject any proposal with probability close to 1; these algorithms will neither add or remove virtual points. The choice of λ is therefore an interesting problem. In Rao and Teh (2013), an appropriate choice of λ is made and only virtual points are added or removed.

2.4 Piecewise homogeneous processes

Our applications to CTBNs require a generalisation of the algorithms in previous subsections so that they can treat a Markov process which 'switches from one regime to another'; the process itself is not time homogeneous, but it is piecewise time homogeneous. Let $t_{\min} < r_1 < \ldots < r_k < t_{\max}$ be points which partition the interval $[t_{\min}, t_{\max}]$ into sub-intervals $[r_{j-1}, r_j]$, $j = 1, \ldots, k, k+1$, with $r_0 = t_{\min}$ and $r_{k+1} = t_{\max}$. Assume that $(X(t), t_{\min} \le t \le t_{\max})$ is a Markov process with finite state space $\mathcal{X}$, such that $(X(t), r_{j-1} \le t \le r_j)$ is a *homogeneous* Markov process with intensity matrix

Q_j. The end value of X on $[r_{j-1}, r_j]$ is the initial value of X on $[r_j, r_{j+1}]$; $X(r_j-) = X(r_j+) = X(r_j)$. Let ν be the probability distribution of $X(0)$. Below we describe the *redundant representation*, also known as the *uniformization* construction of such chains.

Begin with choosing the *redundant intensities of jumps* $\lambda_j \geq \max_x Q_j(x)$, where $Q_j(x) = -Q_j(x,x)$. The potential jump times $t_{\min} < T_1 < \cdots < T_i < \cdots < T_n < t_{\max}$ are then sampled from a piece-wise homogeneous Poisson process with intensity λ_j on $[r_{j-1}, r_j]$. This can be done using the algorithm below.

for $j := 1$ to $k+1$ **do**
  Sample n_j from $Poisson(\lambda_j(r_j - r_{j-1}))$ distribution.
  for $l = 1$ to n_j **do**
    Sample $T_j^{(l)}$ from Uniform(r_{j-1}, r_j).
  end for
end for
Sort all points $T_j^{(l)}$, $l = 1, \ldots, n_j$, $j = 1, \ldots, k+1$ in increasing order and rename them $T_1, \ldots, T_n$.

where $n = \sum_j n_j$. Some of n_js (or even all of them) can be 0. At the second stage, we mark points that have been generated $T_1 = t_1, \ldots, T_n = t_n$ by simulating a *redundant skeleton Markov chain* $X_1, \ldots, X_n$, similarly to the homogeneous case. Here the chain is no longer time homogeneous and its transition probabilities depend on $(t_i)_{i=1}^n$. More precisely, they depend on the intervals to which subsequent t_is belong. Write

$$P_j(x,x') = \begin{cases} Q_j(x,x')/\lambda & \text{if } x \neq x'; \\ 1 - Q_j(x)/\lambda & \text{if } x = x' \end{cases}$$

and let

$$\mathbb{P}[X_i = x' | X_{i-1} = x] = P_j(x,x') \tag{6}$$

whenever $t_i \in [r_{j-1}, r_j)$.

The rest of the construction is exactly the same as for time homogeneous processes. Let $X(t) = X_{i-1}$ for $t_{i-1} \leq t < t_i$, ignoring the 'change of regime' points r_j. Nota bene: we tacitly assumed that these points r_j are fixed. In our applications to CTBNs, the change of regime at a node occurs if some parent node changes state, so the r_js are random. This does not present any difficulty, since the whole construction may be applied conditionally.

The proposal moves and accompanying acceptance rules described in the previous subsections are easy to modify for piecewise homogeneous process. We briefly describe the necessary modifications, omitting the details, which are rather self-evident. In the sequel, write $r(t) = j \quad \text{if } t \in [r_{j-1}, r_j)$.

- *ChangeTime* needs no modifications if it is applied *separately* to any interval of homogeneity $[r_{j-1}, r_j]$. This means that r_{j-1} and r_j take over the role of the endpoints 0 and 1, respectively and we move points t_i belonging to (r_{j-1}, r_j). Thus a jump time t_i never moves from one prior regime to another.
- Alternatively, *ChangeTime* can be applied *globally* to the whole interval $[t_{\min}, t_{\max}]$ on the condition that we use 'uniform uniformization' on this interval; that is, we choose $\lambda_j = \lambda \geq \max_j \max Q_j(x)$ then a jump time t_i can be moved to a different interval of homogeneity.
- In *ChangeState* we use the skeleton transition probabilities linked to the jump times instead of a single P. More precisely, when updating x_i to x_i', we sample the new state x_i' with probability proportional to $P_{r(t_i)}(x_{i-1}, x_i')P_{r(t_{i+1})}(x_i', x_{i+1})$ if $0 < i < n$. Similarly, x_0' is sampled with probability proportional to $\nu(x_0')P_{r(t_1)}(x_0', x_1)$ and x_n' with probability proportional to $P_{r(t_n)}(x_{n-1}, x_n')$.
- In *AddRandomPoint* and *EraseRandomPoint* we modify the acceptance probabilities analogously.

3 Continuous time Bayesian networks

First we recall the definition of a CTBN and basic facts about this notion.

3.1 Definitions and notations

Let $(V, \mathcal{E})$ be a directed graph with possible cycles. We write $v \to w$ instead of $(v,w) \in \mathcal{E}$, whenever the graph is fixed. For $v \in V$ let $\mathrm{pa}(v) = \{w : w \to v\}$ be the set of parents of v. Suppose $\mathcal{A}_v$ is the alphabet of possible states of node v. We consider a class of continuous time stochastic processes on the product space $\mathcal{X} = \prod_{v \in V} \mathcal{A}_v$. Thus a state $x \in \mathcal{X}$ is a configuration $x = (x_v) = (x_v)_{v \in V}$, where $x_v \in \mathcal{A}_v$. As usual, if $W \subseteq V$ then we write $x_W = (x_v)_{v \in W}$ for the configuration x restricted to nodes in W. We also use the notation $\mathcal{X}_W = \prod_{v \in W} \mathcal{A}_v$, so that $x_W \in \mathcal{X}_W$. The set $W \setminus \{v\}$ will be denoted by $W - v$ and $V \setminus \{v\}$ simply by $-v$. Suppose we have a family of functions $Q_v : \mathcal{X}_{\mathrm{pa}(v)} \times (\mathcal{A}_v \times \mathcal{A}_v) \to [0, \infty)$. For fixed $c \in \mathcal{X}_{\mathrm{pa}(v)}$, we consider $Q_v(c; \cdot, \cdot)$ as a conditional intensity matrix (CIM) at node v (only the off-diagonal elements of this matrix have to be specified since those on the diagonal are irrelevant). The state of a CTBN at time t is a random element $X(t)$ of the space $\mathcal{X}$ of configurations. Let $X_v(t)$ denote its vth coordinate. The process $\boldsymbol{\Xi} = ((X_v(t))_{v \in V}, t \geq 0)$ is assumed to be Markov and its evolution can be described informally as follows. The

transition intensities at node v depend on the current configuration of the parent nodes. If the parent configuration changes, then node v switches to other transition intensities. If $x_v \neq x'_v$ then

$$\mathbb{P}\left[X_v(t+\mathrm{d}t) = x'_v | X_{-v}(t) = x_{-v}, X_v(t) = x_v\right] = Q_v(x_{\mathrm{pa}(v)}; x_v, x'_v)\,\mathrm{d}t.$$

Formally, CTBN is a THMP with transition intensities given by

$$Q(x,x') = \begin{cases} Q_v(x_{\mathrm{pa}(v)}; x_v, x'_v) & \text{if } x_{-v} = x'_{-v} \\ & \text{and } x_v \neq x'_v \text{ for some } v; \\ 0 & \text{if } x_{-v} \neq x'_{-v} \text{ for all } v, \end{cases}$$

for $x \neq x'$ (of course, $Q(x,x)$ must be defined by subtraction in the usual way to ensure that $\sum_{x'} Q(x,x') = 0$).

3.2 Probability densities of CTBNs

An important special case of evidence is the *complete* observation of some nodes of the CTBN. To compute the posterior distribution over unobserved nodes, we need the likelihood; the probability density of the observed trajectories. Formulae for densities of general HMMs can be obtained from (3) by 'intergrating out the virtual jumps'. Such formulae appear in many papers, e.g. Nodelman et al. (2003) and Rao and Teh (2013), equation (2). The latter reference also contains a comprehensive discussion about the reference measure with respect to which the density is computed. This is not important for our purposes. Below, we recall a formula specialized to CTBNs, which is equivalent to expressions given in Nodelman et al. (2003). As before, we set a finite time horizon, say $[t_{\min}, t_{\max}] = [0,1]$, and consider a CTBN process

$$\boldsymbol{\Xi} = ((X_v(t))_{v\in V}, 0 \leq t \leq 1)\,.$$

Recall that the state space is $\mathcal{X} = \prod_v \mathcal{A}_v$ and the transition intensities are described by CIMs Q_v. We need the following notations:

> Let $n_v^{\boldsymbol{\xi}}(c;\ a,a')$ denote the number of jumps from $a \in \mathcal{A}_v$ to $a' \in \mathcal{A}_v$ at node v, which occurred when the parent configuration was $c \in \mathcal{X}_{\mathrm{pa}(v)}$.
> Let $t_v^{\boldsymbol{\xi}}(c;\ a)$ be the length of time that node v was in state $a \in \mathcal{A}_v$ and the parent configuration was $c \in \mathcal{X}_{\mathrm{pa}(v)}$.

The density of the sample path $\boldsymbol{\xi} = ((x_v(t))_{v\in V}, 0 \leq t \leq 1)$ is the following:

$$p(\boldsymbol{\xi}) = \nu\left(x(0)\right) \prod_{v\in V} \varrho(\boldsymbol{\xi}_v \| \boldsymbol{\xi}_{\mathrm{pa}(v)}), \tag{7}$$

where $x(0) = (x_v(0))_{v\in V}$ is the configuration at time 0, ν is the initial distribution and

$$\varrho(\boldsymbol{\xi}_v \| \boldsymbol{\xi}_{\mathrm{pa}(v)}) = \left\{ \prod_{c\in\mathcal{X}_{\mathrm{pa}(v)}} \prod_{a\in\mathcal{A}_v} \prod_{\substack{a'\in\mathcal{A}_v \\ a'\neq a}} Q_v(c;\ a,a')^{n_v^{\boldsymbol{\xi}}(c;\ a,a')} \right\} \left\{ \prod_{c\in\mathcal{X}_{\mathrm{pa}(v)}} \prod_{a\in\mathcal{A}_v} \exp\left[-Q_v(c;\ a) t_v^{\boldsymbol{\xi}}(c;\ a)\right] \right\}, \tag{8}$$

To give a clear and useful interpretation of (8), let us recall the notion of *conditioning by intervention*. This concept is well understood in the context of static BNs and acyclic digraphs (Lauritzen (1996) and the references therein). The idea can easily be carried over to CTBNs with graphs that are possibly cyclic. For $W \subset V$, let $\boldsymbol{\Xi}_W = ((X_v(t))_{v\in W}, 0 \leq t \leq 1)$ be the process restricted to the nodes in W. Write $\mathrm{pa}(V \setminus W) = \{w : w \in W \text{ and } w \to v \text{ for some } v \notin W\}$ and

$$\varrho(\boldsymbol{\xi}_{V\setminus W} \| \boldsymbol{\xi}_W) = \varrho(\boldsymbol{\xi}_{V\setminus W} \| \boldsymbol{\xi}_{\mathrm{pa}(V\setminus W)}) = \prod_{v\in V\setminus W} \varrho(\boldsymbol{\xi}_v \| \boldsymbol{\xi}_{\mathrm{pa}(v)}). \tag{9}$$

Suppose that the trajectory $\boldsymbol{\Xi}_W = \boldsymbol{\xi}_W$ is fixed. Imagine that we remove the arrows of the graph $(V,\mathcal{E})$ which lead into W and allow the nodes outside W to evolve according to the CTBN dynamics, always using the current values of the fixed trajectory $\boldsymbol{\xi}_W$ in the CIMs. Strictly speaking, the resulting stochastic process on $\mathcal{X}_{V\setminus W}$ is a piecewise time homogeneous Markov chain where the intensity matrix is constant on time intervals where the configuration $x_W(t)$ is constant. If we start at a deterministic initial state, say $x_{V\setminus W}(0)$, then the density of the process is proportional to $\varrho(\boldsymbol{\xi}_{V\setminus W} \| \boldsymbol{\xi}_{\mathrm{pa}(V\setminus W)})$. Thus (9) corresponds to the condition-by-intervention transition rule of a CTBN on $\mathcal{X}_{V\setminus W}$, given $\boldsymbol{\xi}_W$.

4 A Gibbs sampler for CTBNs

The main idea is a straightforward application of 'Metropolis within Gibbs'. We embed the algorithms from the previous section into a GS which updates single nodes one after another. In what follows, we focus on a special form of evidence. Assume that $W \subset V$ is the set

of nodes which are observed completely over some period of time, say $[0,1]$. We update a node $v \notin W$ using one or more Metropolis steps which preserve probability distribution

$$p(\boldsymbol{\xi}_v|\boldsymbol{\xi}_{-v}) \propto p(\boldsymbol{\xi}_v, \boldsymbol{\xi}_{-v}) = p(\boldsymbol{\xi}).$$

This general strategy has many variants. We can use either a random scan or a systematic scan Gibbs sampler (GS). The choice of Metropolis moves and the number of such moves for a single Gibbs step can be specified in many ways. Instead of single nodes we can update subsets of $V \setminus W$ simulataneously (block GS). The effeciency of such variants is clearly problem-dependent and will not be discussed here. We will only explain how to apply the Metropolis moves from Section 2 to GS for CTBNs.

To apply functions from Section 2, we need to specify the prior distribution at node v, which is the distribution of a piecewise homogeneous Markov process with state space $\mathcal{A}_v$, and the likelihood. In Section 2, we worked with redundant representations $\mathbf{x}$ rather than sample paths $\boldsymbol{\xi}$. This causes no problem, since any move that preserves the probability distribution of an MPP also preserves the inherited distribution of the Markov process represented by the MPP. Let us therefore abuse notation and use Greek letters as arguments of π and L. With this convention, we express $p(\boldsymbol{\xi}) \propto \pi(\boldsymbol{\xi}_v)L(\boldsymbol{\xi}_{-v}|\boldsymbol{\xi}_v)$. To implement efficient algorithms over a CTBN, we have to assume some special stucture of the initial distribution ν. We sketch two scenarios below.

In some applications we may assume, as in Nodelman et al. (2002) and Niemiro (2014), that the initial distribution is specified as a static Bayesian network. Suppose that $(V, \mathcal{E}_0)$ is a *directed acyclic graph* (DAG) and that the distribution ν factorizes as:

$$\nu\big((x_v(0))_{v\in V}\big) = \prod_v \nu(x_v(0)|x_{\mathrm{pa}^0(v)}(0)), \tag{10}$$

where $\mathrm{pa}^0(v)$ refers to the set of parents with respect to $\mathcal{E}_0$. The condition-by-intervention initial distribution is then defined in the standard way:

$$\nu\big(x_{V\setminus W}(0)\|x_W(0)\big) = \prod_{v\notin W} \nu(x_v(0)|x_{\mathrm{pa}^0(v)}(0)).$$

If (10) holds, we may choose

$$\begin{aligned}\pi(\boldsymbol{\xi}_v) = p(\boldsymbol{\xi}_v\|\boldsymbol{\xi}_{-v}) &= \nu(x_v(0)\|x_{-v}(0))\varrho(\boldsymbol{\xi}_v\|\boldsymbol{\xi}_{-v}),\\ L(\boldsymbol{\xi}_{-v}|\boldsymbol{\xi}_v) = p(\boldsymbol{\xi}_{-v}\|\boldsymbol{\xi}_v) &= \nu(x_{-v}(0)\|x_v(0))\varrho(\boldsymbol{\xi}_{-v}\|\boldsymbol{\xi}_v).\end{aligned}$$

Of course, $\varrho(\boldsymbol{\xi}_v\|\boldsymbol{\xi}_{-v})$ depends on $\boldsymbol{\xi}_{-v}$ only through $\boldsymbol{\xi}_{\mathrm{pa}(v)}$ and $\varrho(\boldsymbol{\xi}_{-v}\|\boldsymbol{\xi}_v)$ depends on $\boldsymbol{\xi}_{-v}$ only through $\boldsymbol{\xi}_{\mathrm{ch}(v)}$. Similarly, the ν terms depend only on the variable / parent configurations $x_v(0)/x_{\mathrm{pa}^0(v)}(0)$ where pa^0 denotes the parent set with respect to the DAG $\mathcal{E}_0$. The likelihood is computed according to formulae in Section 3. It follows that the functions of Section 2 can be implemented efficiently.

Another possible scenario is that we are able to evaluate the full conditional distributions at time 0, $\nu(x_v(0)|x_{-v}(0))$. Then we can write:

$$\begin{aligned}\pi(\boldsymbol{\xi}_v) &= \nu(x_v(0)|x_{-v}(0))\varrho(\boldsymbol{\xi}_v\|\boldsymbol{\xi}_{-v}),\\ L(\boldsymbol{\xi}_{-v}|\boldsymbol{\xi}_v) &= \nu(x_{-v}(0))\varrho(\boldsymbol{\xi}_{-v}\|\boldsymbol{\xi}_v) \propto \varrho(\boldsymbol{\xi}_{-v}\|\boldsymbol{\xi}_v).\end{aligned}$$

The factor $\nu(x_{-v}(0))$ need not be evaluated.

The Metropolis within GS produces a sequence of processes $\boldsymbol{\Xi}^1_{V\setminus W}, \ldots, \boldsymbol{\Xi}^m_{V\setminus W}, \ldots$ which is a Markov chain with stationary distribution $\pi(\cdot|\boldsymbol{\xi}_W)$. The estimator of the inference probability is

$$\hat{\pi}_m(\cdot|\boldsymbol{\xi}_W) = \frac{1}{m}\sum_{j=1}^{m} \mathbb{I}\left\{\boldsymbol{\Xi}^j_{V\setminus W} \in \cdot\right\}. \tag{11}$$

The funcions of Section 2 which change the number of jumps ensure that the chain is irreducible and thus ergodic. Estimator (11) is strongly consistent.

5 Simulation results

We applied our algorithm to a simple network. Consider a CTBN with two vertices X and Y, each of them binary, i.e. the alphabet of states is $\mathcal{A}_X = \mathcal{A}_Y = \{1,2\}$. The process at node X is hidden, node Y is fully observable. The graph is

$$X \longrightarrow Y$$

The parameters of the process are

- the transition intensity matrix $Q_X(x,x')$;
- the conditional intensity matrices $Q_Y(x;y,y')$, where x denotes the state of the parent node X.

We experimented with two sets of parameters. In Example 1, the current state of X affects the probabilities with which Y 'chooses its state'. In Example 2, the current state of X affects the frequency of jumps of Y. To give the algorithm a fighting chance to restore the path of X given information about Y, we have to choose transition rates of Y substantially higher than that of X. In both exapmples we used the LW algorithm and our Metropolis-type algorithm. In the latter, we iterated functions *ChangeTime*, *ChangeState*, *AddRandomPoint* and *EraseRandomPoint* cyclically.

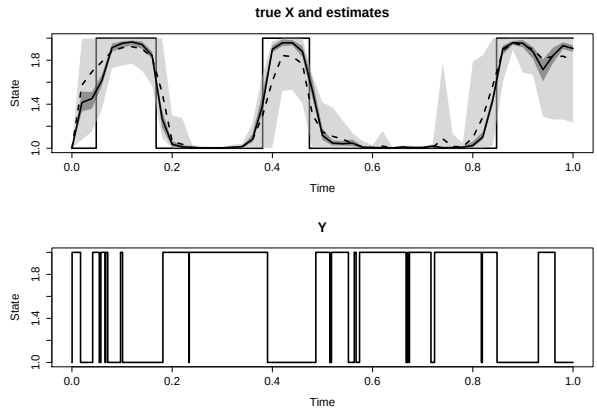

Fig. 1 Results of Estimation (above), Sample paths of X and Y (below) for Example 1. Above: Dashed - likelihood weighting, solid- MCMC, lighter shadow - variability of likelihood weighting, darker shadow - variability of MCMC

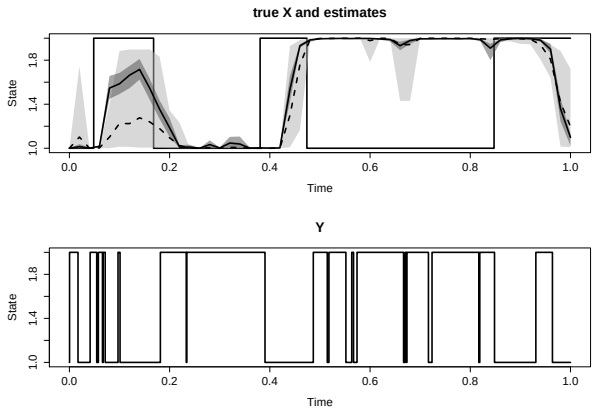

Fig. 2 Results of Estimation (above), Sample paths of X and Y (below) for Example 2. Above: Dashed - likelihood weighting, solid- MCMC, lighter shadow - variability of likelihood weighting, darker shadow - variability of MCMC

Examples 1 and 2

Consider a THMP (X, Y) with state space $\{1, 2\}^2$, where X itself is a THMP with transition matrix

$$Q_X = \begin{pmatrix} -4 & 4 \\ 5 & -5 \end{pmatrix}$$

and the transition matrix for Y depends on X. Y is observed and we want to make inferences about the hidden Markov chain X. For Example 1, we take:

$$Q_{Y|X=1} = \begin{pmatrix} -100 & 100 \\ 20 & -20 \end{pmatrix}, \qquad Q_{Y|X=2} = \begin{pmatrix} -20 & 20 \\ 100 & -100 \end{pmatrix}$$

so that for Y has a substantially greater probability of being in state 2 when $X = 1$ than when $X = 2$ and has substantially greater probability of being in state 1 when $X = 2$ than when $X = 1$. This is illustrated by a sample path of the process (X, Y) shown in Fig. 1 The time interval under consideration is always $[t_{\min}, t_{\max}] = [0, 1]$.

For Example 2, we take the same Q_X, but with $Q_{Y|X}$ given below.

$$Q_{Y|X=1} = \begin{pmatrix} -100 & 100 \\ 100 & -100 \end{pmatrix}, \qquad Q_{Y|X=2} = \begin{pmatrix} -2 & 2 \\ 2 & -2 \end{pmatrix}$$

The jump frequency for Y is substantially higher for $X = 1$ than for $X = 2$, but the jumps from $Y = 1$ to $Y = 2$ and from $Y = 2$ to $Y = 1$ are equally likely. This is illustrated by the sample paths of the process shown in Fig. 2 The type of information transmitted from X to Y is clearly quite different in these two examples.

Fig. 1 shows the estimation results for Example 1, while Fig. 2 shows the estimation results for Example 2. More precisely, these figures depict the posterior probability of $X(t) = x$ given the whole path $(Y(t), 0 \leq t \leq 1)$. The results obtained by LW are given by dotted lines, while the results of our MCMC algorithm are given by dashed lines. The LW estimator and our MCMC algorithm given by (11) are applied to $\hat{\pi}_m(X(t) = 1|\boldsymbol{Y} = \boldsymbol{\eta})$ for a grid of t-points, where $\boldsymbol{Y} = \boldsymbol{\eta}$ is the observed path of Y. The solid broken line is the true unobserved path of the hidden node X.

This experiment gives a spectacular illustration of the degeneracy of weights phenomenon of the LW algorithm. For a sample of size $m = 10000$, the cumulative sums of the 10 largest weights (normalized to sum to 1) are shown in Table 1 for Example 1 and in Table 2 for Example 2.

0.538	0.906	0.939	0.955	0.967
0.974	0.981	0.984	0.986	0.988

Table 1 Cumulative sums of the 10 largest weights in the LW algorithm for Example 1

0.589	0.741	0.781	0.803	0.825
0.847	0.867	0.886	0.899	0.912

Table 2 Cumulative sums of the 10 largest weights in the LW algorithm for Example 2

Thus, for Example 1, 10 of 10000 points carry about 98.8% of the total mass. Roughly speaking, 9990 sampled points are effectively useless. The results for Example 2 are slightly better; the l0 largest weights account for 91.2% of the value. In a more realistic situation, when the problem is to compute the posterior

over a set of hidden nodes of a complicated CTBN, the phenomenon of degeneracy may become substantially worse.

For both Example 1 and Example 2, the number of steps of the Metropolis-type algorithm is $m = 10000$. Each step consists of consecutive applications of *ChangeTime*, *ChangeState* and then *AddRandomPoint* or *EraseRandomPoint*. Parameter λ is chosen as $2.5 \times \max_x Q_X(x)$. The acceptance rate was approximately 0.5 in the experiment for both examples. The results were similar for both examples, even though the type of evidence was quite different. The results represent a substantial improvement over those for LW.

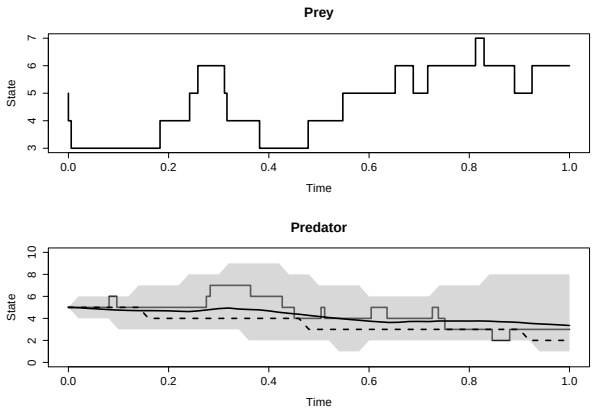

Fig. 3 Results for Lotka Volterra example: dashed line - posterior median, solid - posterior mean, shadow - 0.1 and 0.9 posterior quantiles

Example 3: A Lotka Volterra Model

Let (X, Y) be a Markov chain with intensities:

$$
\begin{aligned}
Q_{(x,y)\to(x,y+1)} &= y(\alpha - \beta y) \vee 0 \,, \\
Q_{(x,y)\to(x,y-1)} &= \gamma y x \wedge M \,, \\
Q_{(x,y)\to(x,z)} &= 0 \text{ for } z \neq y \pm 1, \\
Q_{(x,y)\to(x+1,y)} &= \delta x y \wedge M \,, \\
Q_{(x,y)\to(x-1,y)} &= \eta x \wedge M \,, \\
Q_{(x,y)\to(z,y)} &= 0 \quad \text{for } z \neq x \pm 1,
\end{aligned}
$$

where $\alpha, \beta, \gamma, \delta, \eta$ are non-negative constants and M is a truncation parameter so that uniformization may be used. In this model, Y is the *prey*, while X is the *predator*. The prey is *hidden*, while the predator is *observed*. In the absence of predation, the prey population grows, with limitations on the growth rate due to the carrying capacity of the environment. At the same time, each prey is killed off with intensity proportional to the number of predators. In the absence of any prey, the predator's death rate is exponential, while the prey contributes to the predator's growth rate. The coefficients are chosen so that the Markov chain (X, Y) has a stationary distribution. Such models are discussed in Murray (1993). The results are illustrated in Fig. 3.

This example illustrates that, although the intensities need to remain bounded so that uniformization may be used, we do not need a bounded state space.

6 Discussion

We have presented a new Metropolis-Hastings MCMC algorithm for detecting hidden variables in a CTBN. The algorithm presented here has some similarities to that of Rao and Teh (2013), but operates on essentially different principles. The two algorithms are suited to different situations. Our algorithm is substantially more local in flavour. Firstly, in contrast to Rao and Teh, each move of *Add/EraseRandomPoint* in our algorithm only needs to update the sufficient statistics for three points, while a single move of the Rao and Teh algorithm re-evaluates the entire trajectory. The Rao-Teh algorithm uses an FFBS approach which requires, at each step the multiplication of transition matrices, with a cost of $\mathcal{O}(n|\mathcal{X}|)$ where $|\mathcal{X}|$ is the number of elements in the state space and n is the number of jumps, including *virtual* jumps. Broadly speaking, the Rao-Teh algorithm outperforms our algorithm in situations with a state space of moderate size, but the cost of the Rao - Teh algorithm increases linearly with the size of the state space and it cannot perform reliably with *infinite* state space; it is well known that the stationary distribution for a truncated problem may be substantially different from the target stationary distribution; the Rao-Teh algorithm may not be able to detect this. The cost of our algorithm is broadly independent of the size of the state space. The Lotka-Volterra example illustrates the performance of our algorithm in a situation where the state space is unbounded and indicates that it can give a satisfactory performance in such situations.

An interesting problem is the choice of $\lambda > \max_x Q(x)$. This is also an issue for Rao and Teh, but in their situation the answer is reasonably clear cut. For their algorithm, large values of λ increase mobility and therefore efficiency, but at the same time increase cost, because the number of virtual jumps increases. The value of λ is therefore the largest permitted by constraints of cost.

In the situation here, the choice $\lambda = \max_x Q(x)$ can lead to unfortunately low acceptance probabilities; if (for example) point x_i is proposed via $P(x_{i-1}, .)$, but $x_i = x_{i+1}$, then for $\max_x Q(x) = Q(x_i)$, $P(x_i, x_{i+1}) = 0$, so that $a(\mathbf{x}', \mathbf{x}) = 0$. On the other hand, if λ is too

large, the skeleton will have more points, many of them virtual, which decreases efficiency.

In conclusion, the algorithm presented here gives a new contribution, which complements existing approaches; some of the important applications are not within the scope of existing algorithms.

The algorithm can also be extended in a straightforward manner to the situation where instead of the whole trajectory, the observed data is the trajectory sampled at a finite number of fixed time points.

References

Paul Dagum and Michael Luby. An optimal approximation algorithm for bayesian inference. *Artificial Intelligence*, 93(1–2):1–27, 1997.

Yu Fan, Jing Xu, and Christian R. Shelton. Importance sampling for continuous time Bayesian networks. *Journal of Machine Learning Research*, 11 (Aug):2115–2140, 2010.

P. J. Green. Reversible jump markov chain monte carlo computation and bayesian model determination. *Biometrika*, 82(4):711–732, 1995.

W. K. Hastings. Monte carlo sampling methods using markov chains and their applications. *Biometrika*, 57 (1):97–109, 1970.

Asger Hobolth and Eric A. Stone. Simulation from endpoint-conditioned, continuous-time markov chains on a finite state space, with applications to molecular evolution. *The Annals of Applied Statistics*, 3(3):1204–1231, 2009.

S.L. Lauritzen. *Graphical Models*. Clarendon Press, 1996. ISBN 9780191591228.

Nicholas Metropolis, Arianna W. Rosenbluth, Marshall N. Rosenbluth, Augusta H. Teller, and Edward Teller. Equation of State Calculations by Fast Computing Machines. *The Journal of Chemical Physics*, 21(6):1087–1092, 1953.

J.D. Murray. *Mathematical Biology*. Springer, 1993. ISBN 0-540-57204.

Brenda Ng, Avi Pfeffer, and Richard Dearden. Continuous time particle filtering. In *Proceedings of the 19th International Joint Conference on Artificial Intelligence*, IJCAI'05, pages 1360–1365, San Francisco, CA, USA, 2005. Morgan Kaufmann Publishers Inc.

Wojciech Niemiro. Metropolis algorithm and likelihood weighting for CTBNs. 2014.

U. Nodelman. *Continuous Time Bayesian Networks*. PhD thesis, Department of Computer Science, Stanford University, 2007.

U. Nodelman, C.R. Shelton, and D. Koller. Continuous time Bayesian networks. In *Proceedings of the Eighteenth Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 378–387, 2002.

U. Nodelman, C.R. Shelton, and D. Koller. Learning continuous time bayesian networks. In *Proc. Nineteenth Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 451–458, 2003.

Judea Pearl. Causal diagrams for empirical research. *Biometrika*, 82(4):669–710, 1995. With discussion and a rejoinder by the author.

V. Rao and Y. W. Teh. Fast MCMC sampling for Markov jump processes and extensions. *Journal of Machine Learning Research*, 14:3207–3232, 2013.

Tore Schweder. Composable markov processes. *Journal of applied probability*, 7(2):400–410, 1970.